arXiv:1610.09320v3 [cs.DS] 15 Dec 2016

A Polynomial-time Algorithm for Detecting the Possibility of Braess Paradox in Directed Graphs

Pietro Cenciarelli Daniele Gorla Ivano Salvo
Sapienza University of Rome, Department of Computer Science
cencia@di.uniroma1.it, gorla@di.uniroma1.it, salvo@di.uniroma1.it

Abstract

Vulnerable graphs are those directed multigraphs that can generate *Braess paradox* in Traffic Networks. In this paper, we give a graph-theoretic characterisation of vulnerable directed multigraphs; analogous results appeared in the literature only for undirected multigraphs and for a specific family of directed multigraphs. The proof of our characterisation also provides an algorithm that checks if a multigraph is vulnerable in $O(|V| \cdot |E|^2)$; this is the first polynomial time algorithm that checks vulnerability for general directed multigraphs. The resulting algorithm also contributes to another well known problem, i.e. the directed subgraph homeomorphism problem without node mapping, by providing another pattern graph for which a polynomial time algorithm exists.

1 Introduction

Traffic Networks [1, 2] provide a model for studying selfish routing: non-cooperative agents travel from a source node s to a destination node t. Since the cost (or latency) experienced by an agent while travelling along a path depends on network congestion (and hence on routes chosen by other agents), traffic in a network stabilizes to the equilibrium of a non-cooperative game, where all agents experience the same latency. This phenomenon has been defined by Wardrop [23] in the context of transport analysis.

For example, consider Fig. 1(a), that depicts the so called *Wheatstone network*. In the model, edges are labeled by a function (here, the constants 1 and 0, and the identity function x) that specifies the latency of each edge in terms of the flow it experiences. So, for example, the edge of latency 0 is "ideal", in the sense that, independently on how much traffic travels along it, the passage from u to v is instantaneous. By contrast, the delay on passing from s to u and from v to t is linear in the amount of flow traveling along such edges. Thus, every small autonomous flow particle ϵ will try to use the edge $u \to v$, because in this way it will experience a delay of 2ϵ (instead of $1 + \epsilon$, if it had avoided that edge). Consequently, a Wardrop flow of value 1 assigns all the flow to the path $s\,u\,v\,t$ in the picture; thus, the overall latency is 2.

In traffic networks, a well known and counterintuitive phenomenon is *Braess paradox* [1, 3], that originates when latency at Wardrop equilibrium decreases because of removing edges. The Wheatstone network is a minimal example of Braess paradox: Fig. 1(b) shows

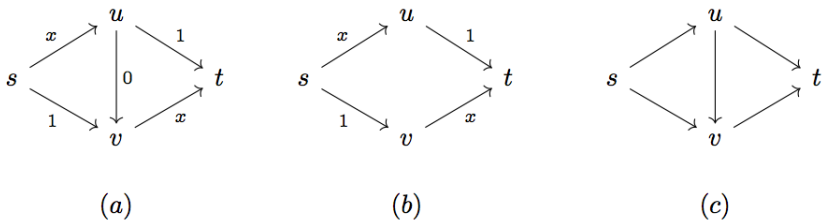

Figure 1: (a) The Wheatstone network; (b) Its optimal subgraph; (c) The graph W.

its optimal subnet. There, a Wardrop flow of value 1 assigns 1/2 to both paths in the network ($s\,u\,t$ and $s\,v\,t$), thus obtaining a latency of 3/2.

Braess paradox has been studied for decades. The results that are most strictly related to ours start with [19] where it is shown that, given a multigraph, a latency function on its edges and the total amount of flow, it is NP-hard to prove whether the resulting net suffers from the Braess paradox or not. An intriguing question raised in [19] (Open Question 1 in Sect. 6.1) is to study Braess paradox from a graph-theoretical perspective, i.e. by considering only a graph and studying whether it admits instances that generate the paradox. This property of a graph has been called *vulnerability*.

A characterisation of vulnerable undirected multigraphs is presented in [16], where it is proved that an undirected graph is vulnerable if and only if it is series-parallel [18]. This characterisation only holds for graphs where every node and every edge lie on at least one simple *st*-path. Later, in [4], the same characterisation is proved for directed multigraphs that satisfy the very same condition, called *irredundancy* therein. However, while checking redundancy and calculating the maximal irredundant subgraph can be efficiently done in undirected graphs [5], the same does not hold for directed graphs. Indeed, in [5] the authors hint at the difficulty of this problem and note that their solution is tractable only for planar digraphs. They also conjecture that the problem of recognizing vulnerable digraphs is untractable, in general.

Contribution In this paper, we disprove this conjecture. First of all, we prove NP-hardness of checking and removing redundancy in directed multigraphs, thus formalizing the informal claim in [5]. Then, we provide a graph-theoretic characterisation of vulnerable directed nets, by proving that a directed multigraph is vulnerable if and only if it contains the graph W underlying the Wheatstone network and depicted in Fig. 1(c). Our constructive proof provides an algorithm to check vulnerability, with an execution time that is $O(|V|\cdot|E|^2)$. This is the first polynomial algorithm we are aware of for checking vulnerability of general directed multigraphs and fixes the issue left open in [5].

Finally, our characterisation and the resulting polynomial time algorithm contribute to another well known problem: the directed subgraph homeomorphism problem without node mapping [8]. This problem consists in fixing a pattern graph (W in our case) and try to find an homeomorphic copy of it within an input graph. In *loc. cit.*, it is suggested

that the general problem is NP-hard. However, in the special case when the pattern graph has all nodes with indegree at most 1 and outdegree at most 2, or indegree at most 2 and outdegree at most 1 (as for W), this is no more necessary, since there exist pattern graphs of this kind for which a polynomial time algorithm exists [10]. However, as far as we know, no polynomial time algorithm for all this class of pattern graphs has been devised so far. Thus, our work also contributes to this research line, by providing another pattern graph for which a polynomial time algorithm exists.

Related work As we already said, the most strictly related papers are [4, 5, 16].

In [16], vulnerable undirected multigraphs are characterised as the non-series-parallel ones that, in turn, are those not containing W as homeomorphic subgraph. Moreover, this work also characterises the undirected networks where all Wardrop equilibria are weakly Pareto efficient; these turn out to be the nets with linearly independent routes (i.e., those nets where every st-path has at least one edge that belongs only to that path). Linear independence of routes is also shown to be the characterisation of those undirected nets that are efficient under heterogeneous players (i.e., nets where different players can have different latencies on the same edges). Furthermore, linear independence is characterised by not having as homeomorphic subgraph any of three elementary nets (one of which is W).

In [4, 5], directed networks are considered, both in their single commodity version (i.e., with just one pair of source and target) and in the multicommodity one (i.e., with many such pairs). Assuming irredundancy of the net (i.e., that every node and edge lies on at least one st-path, for some st pair), they prove that: (1) for single commodity nets, vulnerability coincides with not being series-parallel (like in [16]); and (2) for multicommodity nets, vulnerability coincides with not being series-parallel for at least one maximal st-irredundant subnet, for some st pair. As we show in this paper, checking irredundancy of a net and calculating the maximal irredundant subnet are computationally untractable problems. Indeed, in [5] an algorithm for checking vulnerability is provided only for undirected and for planar directed nets.

An orthogonal bunch of works [13, 19, 20] has been devoted to the complexity of estimating the Braess ratio (i.e., the maximum ratio between the Wardrop latency of G and of any its subnet) and the price of anarchy (i.e., the worst-case ratio between the values of any Wardrop flow and of the optimal one). These works show that both measures have very strong inaproximability, both for the single and for the multicommodity directed scenarios.

Other works have been carried out to graph-theoretically characterise networks with similar kinds of games [7, 11, 12, 15, 17]. These works differ from ours in the kind of efficiency one aims at or in the model of the game. For example, in [15] the aim is to have all players with the same latency in all equilibria, whereas in [7] the aim is to characterise nets whose equilibria all minimize the maximum latency of every path, in a fremework where every player can only choose one path and send along it just one information unit.

Finally, an orthogonal paper is [9], where Braess paradox is generalized to all congestion games; structures that do not suffer of Braess paradox are then characterised in terms of matroids. Their elegant result, differently from ours, it is not directly related to graph-theoretic concepts, nor it provides any algorithmic procedure.

Organization of the paper We start in Section 2 by giving the basic notions on traffic networks and Braess paradox; this leads to the definition of vulnerability. Then, in Section 3 we show that the only existing characterisation of vulnerability for directed multigraphs [4, 5] cannot yield a polynomial algorithm for general multi-digraphs, unless P = NP. In Section 4 we provide our characterisation. In Section 5 we show how this characterisation yields a polynomial time algorithm and then show how this can be used for the directed subgraph homeomorphism problem. In Section 6 we conclude the paper.

2 Traffic Networks, Braess Paradox and Vulnerability

We start by providing the necessary background, by essentially following the presentation in [19]. A *directed multigraph* (or *multi-digraph*) $G = (V, E)$ consists of a set V of *vertices* (or *nodes*) and a set E of *edges*. Every edge relates a pair of vertices; if e relates (u, v), we say that e is an *output* of u and an *input* of v. We denote with $in(x)$ and $out(x)$ the edges entering into x and exiting from x, respectively. When the name of an edge is not relevant but only its extremes are, we denote an edge that relates (u, v) as $u \to v$.

A *path* is a sequence $u_1e_1u_2 \dots u_{n-1}e_{n-1}u_n$ (for $n \geq 1$) of nodes and edges such that e_i relates (u_i, u_{i+1}), for all $i < n$. When only the extremal nodes of the path p are relevant, we shall write $u_1 \overset{p}{\rightsquigarrow} u_n$ (or simply $u_1 \rightsquigarrow u_n$ if we ignore the name given to the path). If each u_i appears just once in p, we say that p is *simple* (or *acyclic*). We fix a *source* node, written s, and a *target* node, written t. The resulting triple (G, s, t) will be called *net* and we shall sometimes write it without specifying the source and target, if clear from the context. An *st-path* is a path from s to t. The set of st-paths in G is denoted by $P(G)$; the set of simple st-paths is denoted by $SP(G)$. We say that a net is *st-connected* if all vertices of G belong to an st-path.

A *flow* for a net (G, s, t) is a function $\varphi : SP(G) \to \mathbb{R}^+$. The value of a flow is the sum of the values sent over all paths. A flow induces a unique flow $\varphi(e)$ on edges: for any edge $e \in E$, $\varphi(e) = \sum_{p \in SP(G): e \in p} \varphi(p)$. A *latency function* $l_e : \mathbb{R}^+ \to \mathbb{R}^+$ assigns to each edge e a latency that depends on the flow on it; as usual, we only consider continuous and non-decreasing latency functions. The latency of an st-path p under a flow φ is the sum of the latencies of all edges in the path under φ, i.e., $l_p(\varphi) = \sum_{e \in p} l_e(\varphi(e))$. If H is a subgraph of G, we denote with $l|_H$ the restriction of the latency function l on the edges of H.

Given a net G, a real number $r \in \mathbb{R}^+$ and a latency function l, we call the triple (G, r, l) an *instance*. A flow φ is *feasible* for (G, r, l) if the value of φ is r. Notice that, since we do not have any constraint on edges or vertices, every r admits at least one feasible flow.

A feasible flow φ for (G, r, l) is at *Wardrop equilibrium* (or is a *Wardrop flow*) if, for all pairs p, q of st-paths such that $\varphi(p) > 0$ and $\varphi(q) > 0$, we have $l_p(\varphi) \leq l_q(\varphi)$. In particular, this implies that, if φ is a Wardrop flow, all st-paths to which φ assigns a positive flow have the same latency. It is known [19] that every instance admits a Wardrop flow and that different Wardrop flows for the same instance have the same latency along all st-paths with a positive flow. Thus, we denote with $L(G, r, l)$ the latency of all st-paths with positive flow at Wardrop equilibrium. In the special case where $r = 0$, we let $L(G, r, l)$ be 0.

Braess paradox [1, 3] originates when latency at Wardrop equilibrium decreases because of removing edges (or equivalently, by raising the latency function on edges): an instance

(G, r, l) suffers from Braess paradox if there is a subgraph of G with a lower latency.

In [19] it is shown that, given an instance (G, r, l), it is NP-hard to prove whether it suffers from the Braess paradox or not. An intriguing question raised in [19] is to study the problem from a graph-theoretical perspective, i.e. by having only G and study whether it admits instances that generate the paradox. This leads to the following definition [19].

Definition 2.1. *A multi-digraph G is* vulnerable *if there exist a value r, a latency function l and a subgraph H of G such that $L(G, r, l) > L(H, r, l|_H)$.*

3 On the Complexity of Irredundancy in Directed Nets

All characterisations of vulnerable nets we are aware of are for *irredundant* nets, that are nets containing only irredundant vertices and edges. As defined in [4] (and derived from [16]), a vertex or an edge is *irredundant* if it appears in a simple st-path, and *redundant* otherwise. Redundant vertices and edges can be safely ignored when studying the Braess phenomenon in a multigraph: they will never be touched by a Wardrop flow, since flows send items only on acyclic paths. The characterisation given by [4] provides a polynomial time algorithm for checking vulnerability for irredundant nets, where it suffices to check whether the net is series-parallel (this can be done in linear time, by following [22]). So, also for redundant nets a polynomial time algorithm would exist, if it existed a polynomial reduction from a redundant net to an irredundant subnet with the same simple paths. As we now show, this is not possible unless P = NP. Furthermore, we will show that also verifying if a net is irredundant is an NP-complete problem. Therefore, checking if a net satisfies the preconditions prescribed for using the characterisation (and, consequently, the algorithm) in [5] cannot be done in polynomial time, unless P = NP.

First, we observe that irredundancy can be checked by only focusing on edges.

Proposition 3.1. *If every edge of (G, s, t) is irredundant, then every node is irredundant.*

Proof. Assume that v is redundant: this means that every st-path containing v is not simple. Fix one of these paths p and consider the edges of p that are input and output of v: they cannot be irredundant, otherwise also v would be. Contradiction. □

We start with an easy characterisation of irredundant edges. By exploiting it, we then move to the problems of: (1) checking whether an edge is irredundant; (2) deriving from any net its maximal irredundant subnet with the same acyclic paths; and (3) checking whether a net is irredundant. We will prove that all these three problems are NP-hard.

Lemma 3.2. *An edge $e = u \to v$ is irredundant in the net (G, s, t) if and only if there exist two node-disjoint paths $s \rightsquigarrow u$ and $v \rightsquigarrow t$.*

Proof. The 'if' part is trivial: the st-path $s \rightsquigarrow u\, e\, v \rightsquigarrow t$ is acyclic and contains e. For the 'only if' part, by definition, there exists a simple st-path $v_0\, e_0\, v_1\, e_1 \dots v_n$ (for $v_0 = s$, $v_n = t$ and $n \geq 1$) such that $e = e_i$, for some $i \in \{0, \dots, n-1\}$. Thus, $s \rightsquigarrow u$ is $v_0\, e_0\, v_1 \dots v_i$ and $v \rightsquigarrow t$ is $v_{i+1} \dots v_n$: these paths are node-disjoint because the original path was simple. □

Definition 3.1. Edge-wise st-irredundancy *(denoted as* **EW-st-IRR***) is the problem of deciding whether, given a multi-digraph $G = (V, E)$, two distinct nodes $s, t \in V$ and an edge $e \in E$, the edge e is irredundant in the net (G, s, t).*

We first observe that **EW-*st*-IRR** belongs to NP. By definition, a certificate for the edge e in G to be irredundant is a simple path from s to t containing e.

Proposition 3.3. EW-*st*-IRR *is NP-hard.*

Proof. Thanks to Lemma 3.2, we show that **EW-*st*-IRR** can be easily used to solve **2-DPP** (the *2-disjoint paths problem*) [14]:

> **2-DPP**: given a directed graph $G = (V, E)$ and four distinct vertices x, y, w, z, decide whether there exist $x \rightsquigarrow y$ and $w \rightsquigarrow z$ node-disjoint.

The polynomial reduction is the following. If G contains an edge $y \to w$, it suffices to check whether such an edge is xz-irredundant. Otherwise, we consider the graph $G' = (V, E \uplus \{e\})$, where e is a new edge relating (y, w), and check whether e is xz-irredundant in G'. □

NP-completeness of checking whether an edge is irredundant in a net will now be used to show NP-hardness of the problem of extracting from a given net its maximal equivalent (with respect to vulnerability) irredundant subnet. Indeed, we shall prove that a subnet is equivalent with respect to vulnerability to the original net, provided that they have the same set of simple paths (see Lemma 4.3 later on).

Definition 3.2. *The* maximal irredundant subnet *problem (denoted as* **MIS***) is the problem of calculating, given a net (G, s, t), its maximal irredundant subnet G' such that $SP(G) = SP(G')$*

Lemma 3.4. MIS(G, s, t) *is unique.*

Proof. Redundant vertices and edges are univocally determined, once fixed G, s and t. Moreover, removing a redundant edge does not change the set of simple paths in G. So, it suffices to delete them from G in order to obtain a subnet (G', s, t) that is irredundant but has the same simple st-paths as G. Indeed, in passing from G to G', we have only removed edges and vertices that either do not belong to any st-path (viz., those vertices that are unreachable from s or that cannot reach t, and edges incident to them) or belong only to cyclic st-paths. To show uniqueness, notice that: if we remove other edges/vertices, the resulting net will loose some acyclic st-path; if we remove not all these edges/vertices, the resulting net will still be redundant. □

Proposition 3.5. MIS *is NP-hard.*

Proof. Thanks to Lemma 3.4, **MIS** can be easily used to solve **EW-*st*-IRR**: an edge is st-irredundant for (G, s, t) if and only if it appears in **MIS**(G, s, t). □

Finally, also the problem of checking whether a given net is irredundant is an NP-complete problem. In principle, this problem could be computationally simpler: there could exist some characterisation that provides a polynomial test to check if a net is redundant or not but without giving any hint on which edges are redundant and which ones are not.

Definition 3.3. *The* st-irredundancy *problem (denoted as* ***st*-IRR***) is the problem of deciding whether, given a multi-digraph $G = (V, E)$ and two distinct nodes $s, t \in V$, the net (G, s, t) is irredundant.*

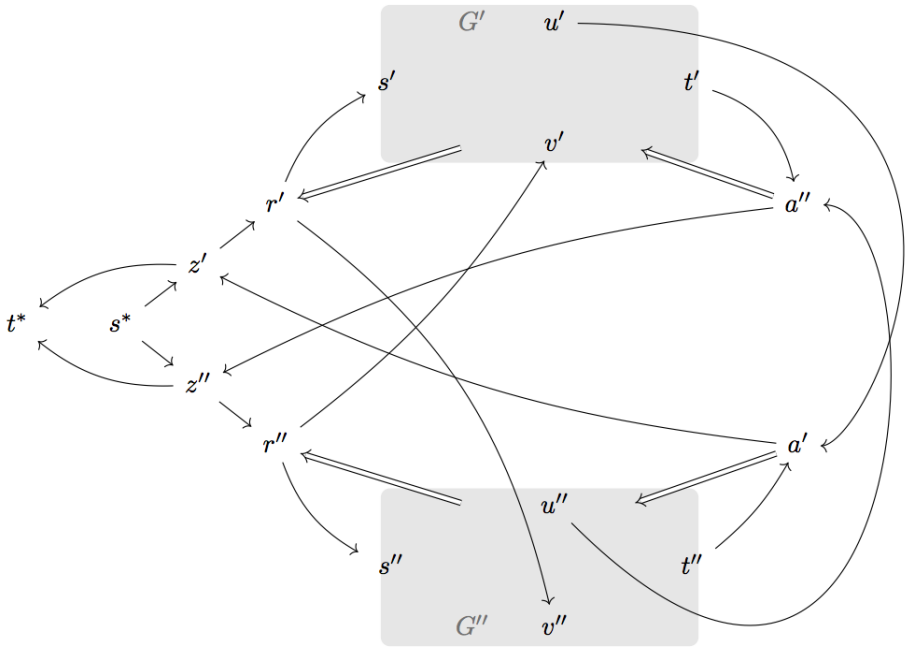

Figure 2: The construction for reducing **EW-*st*-IRR** to ***st*-IRR** (a thick arrow from a node to a grey part representing a graph means $|V|$ edges from the node to every node of the graph; and vice versa)

We first observe that ***st*-IRR** belongs to NP. A certificate for G to be irredundant are $|E|$ simple st-paths such that the i-th path contains the i-th edge of E. We now reduce **EW-*st*-IRR** to ***st*-IRR**: given a net G and an edge $u \to v$, we build a new net G^* such that G^* is redundant if and only if $u \to v$ is redundant in (G, s, t).

Intuitively, our reduction is based on the following ideas: (1) if $u \to v$ is irredundant, then there exists a simple path p in G of the form $s \rightsquigarrow u \to v \rightsquigarrow t$, and (2) for each vertex x, adding new edges of the form $u \to x$ and $x \to v$ does not make irredundant the edge $u \to v$ in G^* if it is redundant in G. By using such additional edges, if $u \to v$ is irredundant in G, we can make irredudant any edge $x \to y$ of G thanks to the path $s \rightsquigarrow u \to x \to y \to v \rightsquigarrow t$ that is simple provided that x and y do not occur in p.

To correctly handle the case in which x or y occur in p, we build a net G^* (with source s^* and target t^*) that consists of two copies of G (call them G' and G'') obtained by decorating each node in G with $'$ and $''$. Each edge $x' \to y'$ in G' (resp. $x'' \to y''$ in G'') will be irredundant in G^* whenever $u \to v$ is irredundant in G thanks to a path of the form $s^* \rightsquigarrow s'' \rightsquigarrow u'' \rightsquigarrow x' \to y' \rightsquigarrow v'' \rightsquigarrow t'' \rightsquigarrow t^*$ (resp. $s^* \rightsquigarrow s' \rightsquigarrow u' \rightsquigarrow x'' \to y'' \rightsquigarrow v' \rightsquigarrow t' \rightsquigarrow t^*$) obtained by adding in G^* paths $u' \rightsquigarrow x''$, $u'' \rightsquigarrow x'$, $x' \rightsquigarrow v''$, and $x'' \rightsquigarrow v'$, for all nodes $x' \in V'$, $x'' \in V''$.

To ensure that redundancy of $u \to v$ in G implies redundancy of $u' \to v'$ and $u'' \to v''$

in G^*, we will define G^* in such a way that a simple path entering in G' via s' (resp. in G'' via s'') can reach the target t^* of G^* only through the node t' (resp. t''), possibly after *only one* detour in G'' (resp. G'). This is guaranteed by additional nodes $z', z'', r', r'', a', a''$ and their incident edges (see Fig. 2). A simple path can enter in G' (resp. G'') only through nodes z' and r' (resp. z'' and r''): this implies that this path can reach t^* only through the path $t' \to a'' \to z'' \to t^*$ (resp. $t'' \to a' \to z' \to t^*$). As a consequence, this path can enter G'' (resp. G') only through $u' \to a' \to x''$ (resp. $u'' \to a'' \to x'$) for some $x' \in V'$ (resp. $x'' \in V''$) and it must eventually come back into G' (resp. G'') through the path $y'' \to r'' \to v'$ (resp. $y' \to r' \to v''$) for some $y' \in V'$ (resp. $y'' \in V''$).

Proposition 3.6. *st*-**IRR** *is NP-hard.*

Proof. Given an instance for **EW-***st*-**IRR** (i.e., a multi-digraph $G = (V, E)$, two distinct vertices $s, t \in V$ and $u \to v \in E$), we first observe that edges entering into s or exiting from t are trivially redundant; so, we can assume that $v \neq s$ and $u \neq t$. Now, we build a new net G^* such that (G^*, s^*, t^*) is irredundant if and only if $u \to v$ is irredundant in (G, s, t). The construction is as follows:

- Create two isomorphic node-disjoint copies of G, call them $G' = (V', E')$ and $G'' = (V'', E'')$, with $V' = \{x' : x \in V\}$, $V'' = \{x'' : x \in V\}$, $E' = \{e' : e \in E\}$ and $E'' = \{e'' : e \in E\}$, where e' and e'' relate (x', y') and (x'', y'') respectively, if e relates (x, y).
- Define G^* as (V^*, E^*), where

$$
\begin{aligned}
V^* &= V' \uplus V'' \uplus \{s^*, z', z'', a', a'', r', r'', t^*\} \\
E^* &= (E' \setminus (in(s') \cup out(t'))) \cup (E'' \setminus (in(s'') \cup out(t''))) \\
&\quad \cup \{s^* \to z', z' \to r', r' \to s', t' \to a'', a'' \to z'', z'' \to t^*, \\
&\qquad s^* \to z'', z'' \to r'', r'' \to s'', t'' \to a', a' \to z', z' \to t^*\} \\
&\quad \cup \{u' \to a', r'' \to v'\} \cup \textstyle\bigcup_{x'' \in V''} \{a' \to x'', x'' \to r''\} \\
&\quad \cup \{u'' \to a'', r' \to v''\} \cup \textstyle\bigcup_{x' \in V'} \{a'' \to x', x' \to r'\}
\end{aligned}
$$

The graphical representation of this costruction is given in Fig. 2.

Irredundancy: Let us first assume that $u \to v$ is irredundant in (G, s, t) and show that every edge in G^* is irredundant. By definition, irredundancy of $u \to v$ implies the existence of a simple st-path in G that contains it; let $s \overset{p_1}{\rightsquigarrow} u \to v \overset{p_2}{\rightsquigarrow} t$ be such a path. First, note that no edge belonging to $in(s) \cup out(t)$ can appear in this path, since all these edges are redundant; for this reason, we have excluded them in G^*.

To show that a generic edge $x' \to y' \in E' \setminus (in(s') \cup out(t'))$ is irredundant, let us consider the path $s^* \to z'' \to r'' \to s'' \overset{p_1''}{\rightsquigarrow} u'' \to a'' \to x' \to y' \to r' \to v'' \overset{p_2''}{\rightsquigarrow} t'' \to a' \to z' \to t^*$, where p_1'' and p_2'' denote p_1 and p_2 within G''. By considering such paths for every $x' \to y' \in E'$, all edges in $\{s^* \to z'', z'' \to r'', r'' \to s'', u'' \to a'', a'' \to x', y' \to r', r' \to v'', t'' \to a', a' \to z', z' \to t^*\}$ are proved irredundant, for every $x', y' \in V'$.

The remaining edges of E^* are "dual" (i.e. have $'$ and $''$ swapped) of the ones considered so far and can be proved irredundant in G^* by considering the dual version (i.e. with $'$ and $''$ swapped everywhere) of the path above.

Redundancy: Let us now assume that $u \to v$ is redundant in (G, s, t) but, by contradiction, assume that there is a simple s^*t^*-path p in G^* that contains $u' \to v'$ (a dual reasoning can be done for proving that also $u'' \to v''$ cannot be irredundant). Let us now yield a contradiction by reasoning on how p has entered into G' for the first time:

1. *Through r':* We first observe that r' can only be reached through z': the only edges different from $z' \to r'$ entering into r' are those of the form $x' \to r'$, for $x' \in V'$; but using these edges would mean that p has already entered into G'. Second, we observe that the only way for reaching t^* with a simple path is through the path $t' \to a'' \to z'' \to t^*$, because z' has already been touched by p. Finally, since G' is isomorphic to G, p cannot lie only within G' because otherwise it would not be simple or would not contain $u' \to v'$. So, it must eventually enter into G'' (and then come back into G').

   We have two possibilities on how p has first entered into G' through r':

   (a) *through the edge $r' \to s'$:* in this case, p cannot reach G'' through a path of the form $x' \to r' \to v''$ otherwise it would be cyclic. Thus, p has reached G'' through a path of the form $u' \to a' \to x''$, for some $x'' \in V''$. However, p cannot pass through $u' \to v'$ before leaving G', otherwise it would be cyclic (since it would touch u' twice). For the same reason, p cannot pass through $u' \to v'$ when it comes back into G'. Contradiction.

   (b) *through the edge $r' \to v''$ and by then entering into G' from G'':* in this case, we have two possible ways for reaching G', i.e. through a path of the form $u'' \to a'' \to x'$ or through a path of the form $x'' \to r'' \to v'$. The first way would force p to touch a'' twice (the second one when passing from t' to t^*). But also the second way would make p cyclic because, to include $u' \to v'$, it should touch v' twice. Contradiction.

2. *Through r'':* In this case, G' cannot be reached through $r'' \to v'$ otherwise, to include $u' \to v'$, p should touch v' twice. Thus, it must be that, from r'', p follows a path of the form $r'' \to s'' \overset{p'}{\rightsquigarrow} u'' \to a'' \to x'$, for some $x' \in V'$. However, having touched a'', the only possible way for p to acyclically reach t^* is by coming back into G'', reaching t'' and then following the path $t'' \to a' \to z' \to t^*$. Thus, p cannot come back into G'' through a path $u' \to a' \to x''$; hence, the only possibility is via a path of the form $y' \to r' \to v''$, for some $y' \in V'$. Moreover, to be acyclic, p has to move from v'' to t'' via a path p'' that only touches vertices that p has not touched before. However, in this way we would have found an $s''t''$- path $s'' \overset{p'}{\rightsquigarrow} u'' \to v'' \overset{p''}{\rightsquigarrow} t''$ that is acyclic and contains the edge $u'' \to v''$. This yields a contradiction because, since G'' is an isomorphic copy of G, every $s''t''$-path in G'' that contains $u'' \to v''$ cannot be simple.

□

4 Characterisation of Vulnerable Multi-Digraphs

A characterisation of vulnerable undirected multigraphs is presented in [16]. In particular, in [16] it is proved that an undirected multigraph is vulnerable if and only if it is series-parallel; moreover, always in [16] it is proved that this holds if and only if the multigraph does not contain a homeomorphic copy of (the undirected version of) W. Our characterisation for multi-digraphs is the combination of these two results (but will be proved directly), thus answering to Open Question 1 in Sect. 6.1 of [19].

To formally state our result, recall [22] that G contains a subgraph homeomorphic to H if H can be obtained from G by a sequence of the following operations:

- remove an edge;
- replace $u \to v$ and $v \to w$ with $u \to w$ and delete vertex v, whenever $in(v) = out(v) = 1$.

In what follows, we assume that G is st-connected. If this is not the case, we can easily find the maximal st-connected subgraph of G (through a visit from s and a backwards visit from t) and check whether the resulting subgraph is vulnerable. Indeed, like for redundancy, edges and vertices that do not lie on any st-path will never play any role in flows and, consequently, cannot influence vulnerability. However, differently from redundancy, the maximal st-connected subnet can be efficiently calculated.

Our main result (Theorem 4.4) states that an st-connected multi-digraph is vulnerable if and only if it contains a subgraph homeomorphic to W. We shall obtain this result in two steps.

First, we characterize vulnerability for acyclic multi-digraphs (Fact 1); this is an immediate corollary of a result in [4], where it is proved that an irredundant multi-digraph is vulnerable if and only if it is not two-terminal series-parallel (TTSP) [18]. By using the standard result from [6], where it is proved that an acyclic directed graph is TTSP if and only if it does not contain a subgraph homeomorphic to W, we can conclude, since acyclic graphs are a special case of irredundant ones.

Fact 1. *Let G be an acyclic st-connected multi-digraph. G is vulnerable if and only if it contains a subgraph homeomorphic to W.*

Second, we reduce the cyclic case to the acyclic one by showing that, in every multi-digraph that does not contain a subgraph homeomorphic to W, we can remove cycles without changing the set of acyclic st-paths (Lemma 4.1); this operation does not affect vulnerability, in the sense that two traffic networks with the same set of acyclic st-paths have exactly the same Wardrop equilibria (Lemma 4.2) and, hence, their underlying graphs are equivalent from the point of view of vulnerability (Lemma 4.3).

Notationally, we write $G' \subseteq G$ whenever $G = (V, E)$ and $G' = (V, E')$, with $E' \subseteq E$; notation $G' \subset G$ has a similar meaning, with $E' \subset E$.

Lemma 4.1. *Let G be an st-connected multi-digraph and C be a simple cycle in G; then, either G contains a subgraph homeomorphic to W or we can remove one edge in C, obtaining a multi-digraph $G' \subset G$ such that $SP(G') = SP(G)$.*

Proof. Without loss of generality, let $v_1, v_2, \ldots, v_n$ (for $n \geq 1$) be the sequence of nodes appearing in C starting from a node v_1 such that there exists a path $s \rightsquigarrow v_1$ that touches C only in v_1. Such a node does exist, since G is st-connected. Then, let v_k be the last node in the sequence $v_1, \ldots, v_n$ such that there exists a path $v_k \rightsquigarrow t$ that touches C only in v_k. Again, such a node does exist, because G is st-connected.

If v_k is the only exit from C, the edge $v_k \to v_{(k \bmod n)+1}$ does not belong to any acyclic st-path in G. Indeed all paths of the form $s \rightsquigarrow v_k \to v_{(k \bmod n)+1} \rightsquigarrow t$ must return in v_k after $v_{(k \bmod n)+1}$ to leave C.

Otherwise, let v_i $(1 \leq i < k)$ be the first node in $v_1, \ldots, v_{k-1}$ that can reach t without passing through other nodes in C and consider nodes $v_{i+1}, \ldots, v_k$. If all paths from s to them pass through other nodes in C, again the edge $v_k \to v_{(k \bmod n)+1}$ does not belong to any acyclic st-path in G. Indeed, we are in the following scenario:

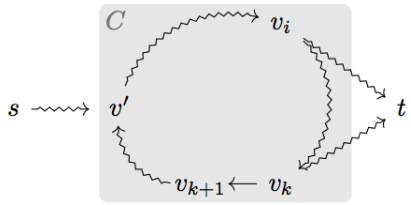

Let us pick up an st-path that uses edge $v_k \to v_{(k \bmod n)+1}$. For what we have just assumed, such a path enters into C in a node $v' \in \{v_{(k \bmod n)+1}, \ldots, v_1, \ldots v_i\}$ and has to pass through the edge $v_k \to v_{(k \bmod n)+1}$. Since all vertices in $v_{(k \bmod n)+1}, \ldots, v_1, \ldots v_i$ can reach t only through v_i or after it, the path is cyclic.

Otherwise, let v_j $(i < j \leq k)$ be such that there exists a path $s \rightsquigarrow v_j$ that does not pass through other nodes in C. In this case, G would contain a subgraph homeomorphic to W, as shown below:

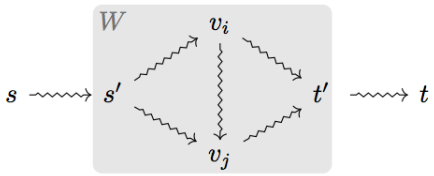

Here, $s \rightsquigarrow s'$ is the longest common prefix of the paths $s \rightsquigarrow v_i$ and $s \rightsquigarrow v_j$ (clearly, we can have $s' = s$); similarly for $t' \rightsquigarrow t$, the longest common suffix of $v_i \rightsquigarrow t$ and $v_j \rightsquigarrow t$. □

Lemma 4.2. *Let $G' \subseteq G$ and φ be a Wardrop flow for (G, r, l). If $SP(G') = SP(G)$ then φ is a Wardrop flow for (G', r, l).*

Proof. By Proposition 2.2 in [19], φ is a Wardrop flow if and only if, for every $p, q \in P(G)$ such that $\varphi(p) > 0$ and $\varphi(q) > 0$, it holds that $l_p(\varphi) \leq l_q(\varphi)$. By definition, φ assigns positive flow only to acyclic st-paths in G; thus, $p, q \in SP(G) = SP(G')$. Moreover, since $G' \subseteq G$, it holds that $P(G') \subseteq P(G)$. Thus, trivially, for every $p, q \in P(G')$ such that $\varphi(p) > 0$ and $\varphi(q) > 0$, it holds that $l_p(\varphi) \leq l_q(\varphi)$. Again by Proposition 2.2 in [19], this means that φ is a Wardrop flow for (G', r, l). □

Lemma 4.3. *Let $G' \subseteq G$. If $SP(G') = SP(G)$ then G is vulnerable if and only if G' is vulnerable.*

Proof. Vulnerability of G' trivially entails vulnerability of G. Let us prove the opposite implication. Let $H \subset G$ be such that $L(H, r, l|_H) < L(G, r, l)$, for some r and l.

If $H \subset G'$, let φ be a Wardrop flow for G. By Lemma 4.2, φ is a Wardrop flow for G' and hence $L(G, r, l) = L(G', r, l)$; thus, $L(H, r, l) < L(G', r, l)$, i.e. G' is vulnerable.

Otherwise, it cannot be $H = G'$, because $SP(H) \subset SP(G) = SP(G')$; indeed, because of Lemma 4.2, if $SP(H) = SP(G)$, we would have $L(H, r, l|_H) = L(G, r, l)$. Thus, there is an edge e belonging to H but not to G'; this means that e only belongs to cyclic st-paths of G, because by hypothesis $SP(G') = SP(G)$. Then, consider H', obtained by removing e from H. Since $P(H) \subseteq P(G)$, we have that e only belongs to cyclic st-paths of H; thus, $SP(H') = SP(H)$. Now, let φ' be a Wardrop flow for H. By Lemma 4.2, φ' is a Wardrop flow for H' and $L(H, r, l|_H) = L(H', r, l|_{H'})$. If $H' \subset G'$, then $L(H', r, l|_{H'}) < L(G', r, l|_{G'})$ and G' is vulnerable. Otherwise, we can find another edge to be removed from H'; but this procedure has to terminate, eventually yielding that G' is vulnerable, as desired. □

Theorem 4.4. *Let G be an st-connected multi-digraph. G is vulnerable if and only if it contains a subgraph homeomorphic to W.*

Proof. For the 'if' part, we know that G admits a subgraph of the form

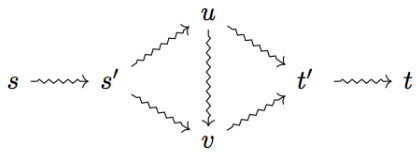

where all paths are node-disjoint, except for the extremal nodes. Let us consider the latency assignment l that assigns:

- 0 to all edges in $u \rightsquigarrow v$, in $s \rightsquigarrow s'$ and in $t' \rightsquigarrow t$;
- x to the first edge in $s' \rightsquigarrow u$ and in $v \rightsquigarrow t'$ and 0 to all the remaining edges in those paths;
- 1 to the first edge in $s' \rightsquigarrow v$ and in $u \rightsquigarrow t'$ and 0 to all the remaining edges in those paths;
- ∞ to all the remaining edges.

Trivially, this reproduces W within G that, consequently, is vulnerable.

For the 'only if' part, if G is not cyclic, the statement follows by Fact 1. If G contains cycles, by contradiction, let us suppose that it does not contain a subgraph homeomorphic to W. Then, by several applications of Lemma 4.1, we can transform G into an acyclic subgraph G' such that $SP(G') = SP(G)$. By Lemma 4.3, since G is vulnerable, G' is vulnerable too; thus, since G' is acyclic, it contains a subgraph homeomorphic to W (by Fact 1); consequently, also G does. Absurd. □

Algorithm 1 Checking Vulnerablity

Input: A directed multigraph $G = (V, E)$, s is the source and t is the target
function *isVulnerable*(G, s, t)
1. Delete from G all vertices not in an st-path and their incident edges
2. **while** G is cyclic **do**
3. Let v_1 be the first vertex touched twice in a DFS visit started from s;
4. Let $C = v_1, \ldots, v_n, v_1$ be the corresponding simple cycle;
5. $S \leftarrow \epsilon$
6. **for** $h = 1, \ldots, n$ **do**
7. **if** t is reachable from v_h in $G \setminus (C \setminus \{v_h\})$ **then**
8. append v_h to the tail of S
9. Let v_i and v_k be the first and the last vertex in S
10. **if** $\exists j$ s.t. $i + 1 \leq j \leq k$ and v_j is reachable from s in $G \setminus (C \setminus \{v_j\})$ **then**
11. **return** TRUE
12. **else**
13. delete $(v_k, v_{(k \bmod n)+1})$ from G;
14. **return** $\neg$TTSP(G)

5 Consequences of the Characterisation

5.1 On Polynomially Checking Vulnerability

Lemma 4.1 is a key result, that also enables for a polynomial time algorithm that checks vulnerability. The algorithm is derived from the proof of that Lemma and it is given in Algorithm 1. There, we used the notation of the proof for naming vertices, so that the reader is guided by the proof in reading the algorithm. Just notice that the case when v_k is the only exit from C is implicitly covered in the algorithm: indeed, in this case $i = k$ and so the guard of the **if** in line 10 is false.

Proposition 5.1. *Checking vulnerability of a net* (G, s, t) *is* $O(|V| \cdot |E|^2)$.

Proof. First, notice that reducing a multi-digraph to its maximum st-connected subgraph (line 1) can be done in $O(|E|)$ by checking reachability from s and backwards reachability from t. Also notice that $TTSP(G)$ is an auxiliary procedure that checks whether G is TTSP or not; it has been devised in [21, 22] and it is linear in the size of the graph.

Concerning the complexity of the algorithm, there are at most $O(|E|)$ iterations of the **while**: indeed, at every iteration either the algorithm terminates finding a subgraph homeomorphic to W (line 11) or one edge is removed from G (line 13). The dominating cost of every iteration is given by the **for** loop at lines 6/8 and by the guard of the **if** at line 10: they cost $O(|V| \cdot |E|)$ since, for every vertex, reachability can be checked via a visit of the resulting multigraph. Overall, the algorithm is $O(|V| \cdot |E|^2)$. □

5.2 On the Directed Subgraph Homeomorphism Problem

Theorem 4.4 and Algorithm 1 have consequences on another, long-standing problem: the directed subgraph homeomorphism [8]. The main problem is the following: fixed a pattern

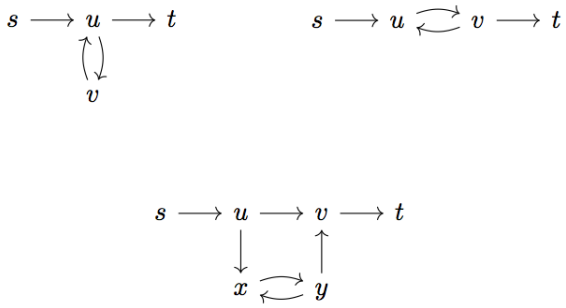

Figure 3: Redundant Graphs that are not vulnerable

graph H, determine whether H is homeomorphic to a subgraph of any given G with respect to a given mapping from the nodes of H to the nodes of G. This problem is shown to be NP-hard, apart from a very simple family of pattern graphs for which a polynomial time algorithm exists. Furthermore, in the conclusions the authors also discuss the problem when the mapping is not given. Our characterisation falls in this latter case.

In [8] it is suggested that the general problem is still NP-hard. However, in the special case when the pattern graph has all nodes with indegree at most 1 and outdegree at most 2, or indegree at most 2 and outdegree at most 1, this is no more necessary, since there exist pattern graphs of this kind for which a polynomial time algorithm exists [10]. However, as far as we know, no polynomial time algorithm for all this class of pattern graphs has been devised so far. Thus, our work also contributes to this research line: W is another pattern graph for which a polynomial time algorithm for the directed subgraph homeomorphism problem without node mapping exists.

Proposition 5.2. *Checking whether a multi-digraph G contains a subgraph homeomorphic to W can be solved in time polynomial to the size of G.*

Proof. Just run Alg. 1 for every pair of distinct vertices in G (ordinately playing the role of the source and target) and return true if and only if at least one of these calls returns true. This algorithm is $O(|V|^3 \cdot |E|^2)$. □

6 Conclusions

We have proved a graph-theoretic characterisation of vulnerable multi-digraphs that generalizes the analogous one for undirected multigraphs [16]. Our characterisation improves the results in [4, 5] (given only for irreduntant directed multigraphs) to general directed multigraphs; indeed, there exist redundant graphs that are not vulnerable, as shown in Fig. 3. Since graphs of this kind can easily appear in real traffic networks, we believe that our characterisation was a necessary step for completing the knowledge of vulnerability.

Also the resulting algorithm was a necessary step: as we proved in this paper, checking irredundancy and calculating the maximal irredundant subnet are NP-hard problems. Of

course, more efficient algorithms for checking vulnerability should be devised to be used in practice, but this was not our aim here.

References

[1] M. Beckmann, C. B. McGuire, and C. B. Winsten. *Studies in the Economics of Transportation*. Yale University Press, 1956.

[2] M. Bell and Y. Iida. *Transportation Network Analysis*. Wiley, 1987.

[3] D. Braess. Über ein paradoxon aus der verkehrsplannung. *Unternehmensforschung*, 12:258–268, 1968.

[4] X. Chen, Z. Diao, and X. Hu. Excluding braess paradox in nonatomic selfish routing. In *Proc. of SAGT15*, volume 9347 of *LNCS*, pages 219–230. Springer, 2015.

[5] X. Chen, Z. Diao, and X. Hu. Network characterizations for excluding braess's paradox. *Theory of Computing Systems*, pages 1–34, 2016.

[6] R. J. Duffin. Topology of series-parallel networks. *Journ. of Math Analysis and Applications*, 10:303–318, 1965.

[7] A. Epstein, M. Feldman, and Y. Mansour. Efficient graph topologies in network routing games. *Games and Economic Behavior*, 66(1):115–125, 2009.

[8] S. Fortune, J. Hopcroft, and J. Wyllie. The directed subgraph homeomorphism problem. *Theoretical Computer Science*, 10(2):111 – 121, 1980.

[9] S. Fujishige, M. Goemans, T. Harks, B. Peis, and R. Zenklusen. Matroids are immune to braess paradox. *Available at: arXiv:1504.07545*, 2015.

[10] M. S. Hecht and J. D. Ullman. Flow graph reducibility. In *Proc. of STOC '72*, pages 238–250. ACM Press, 1972.

[11] R. Holzman and N. Law-yone. Network structure and strong equilibrium in route selection games. *Mathematical Social Sciences*, 46(2):193–205, 2003.

[12] R. Holzman and D. Monderer. Strong equilibrium in network congestion games: increasing versus decreasing costs. *Int. J. Game Theory*, 44(3):647–666, 2015.

[13] H. C. Lin, T. Roughgarden, É. Tardos, and A. Walkover. Stronger bounds on braess's paradox and the maximum latency of selfish routing. *SIAM J. Discrete Math.*, 25(4):1667–1686, 2011.

[14] M. Middendorf and F. Pfeiffer. On the complexity of the disjoint paths problem. *Combinatorica*, 13(1):97–107, 1993.

[15] I. Milchtaich. Topological conditions for uniqueness of equilibrium in networks. *Math. Oper. Res.*, 30(1):225–244, 2005.

[16] I. Milchtaich. Network topology and the efficiency of equilibrium. *Games and Economic Behavior*, 57:321–346, 2006.

[17] I. Milchtaich. Network topology and equilibrium existence in weighted network congestion games. *Int. J. Game Theory*, 44(3):515–541, 2015.

[18] J. Riordan and C. Shannon. The number of two-terminal series-parallel networks. *Journal of Mathematics and Physiscs*, 21:83–93, 1942.

[19] T. Roughgarden. On the severity of Braess's Paradox: Designing networks for selfish users is hard. *J. Comput. Syst. Sci.*, 72(5):922–953, 2006.

[20] T. Roughgarden and É. Tardos. How bad is selfish routing? *Jornal of the ACM*, 49(2):236–259, 2002.

[21] B. Schoenmakers. A new algorithm for the recognition of series parallel graphs. Technical report, CWI - Centrum voor Wiskunde en Informatica, 1995.

[22] J. Valdes, R. Tarjan, and E. Lawler. The recognition of series-parallel digraphs. *SIAM Journal of Computing*, 11:298–313, 1982.

[23] J. Wardrop. Some theoretical aspects of road traffic research. In *Proc. of the Institute of Civil Engineers, Pt. II*, volume 1, pages 325–378, 1952.